Acta Psychologica 81 (1992) 211–241
North-Holland

Dynamic decision making: Human control of complex systems *

Berndt Brehmer
Uppsala University, Uppsala, Sweden

Accepted February 1992

This paper reviews research on dynamic decision making, i.e., decision making under conditions which require a series of decisions, where the decisions are not independent, where the state of the world changes, both autonomously and as a consequence of the decision maker's actions, and where the decisions have to be made in real time. It is difficult to find useful normative theories for these kinds of decisions, and research thus has to focus on descriptive issues. A general approach, based on control theory, is proposed as a means to organize research in the area. An experimental paradigm for the study of dynamic decision making, that of computer simulated microworlds, is discussed, and two approaches using this paradigm are described: the individual differences approach, typical of German work in the tradition of research on complex problem solving, and the experimental approach. In studies following the former approach, the behaviour of groups differing in performance is compared, either with respect to strategies or with respect to performance on psychological tests. The results show that there are wide interindividual differences in performance, but no stable correlations between performance in microworlds and scores on traditional psychological tests have been found. Experimental research studying the effects of system characteristics, such as complexity and feedback delays, on dynamic decision making has shown that decision performance in dynamic tasks is strongly affected by feedback delays and whether or not the decisions have side effects. Although neither approach has led to any well-developed theory of dynamic decision making so far, the results nevertheless indicate that we are now able to produce highly reliable experimental results in the laboratory, results that agree with those found in field studies of dynamic decision making. This shows that an important first step towards a better understanding of these phenomena has been taken.

Our research on dynamic decision making has its origin in analyses of decision making in applied contexts, first in analyses of military decision making and later in analyses of the work of process opera-

* Grants from the Swedish Council of Research in the Humanities and Social Sciences supported the work on this paper as well as the experiments by the author reviewed herein. The author thanks Annica Brehmer, Alex Wearing, and an anonymous reviewer for their helpful comments on an earlier version of this paper.
Correspondence to: B. Brehmer, Dept. of Psychology, University of Uppsala, Box 1854, S-751-48 Uppsala, Sweden.

0001-6918/92/$05.00

tors, in emergency management, and in intensive care. Two things emerged from these analyses.

First, decision making itself was never the primary focus of activity. Instead, decision making was always directed towards some goal. A quote from Klein (personal communication, 1986), describing his work with fire chiefs, summarizes much of our experience, too. When asked about their decisions at the site of the fire, the fire chiefs said: '*We do not make decisions, we fight fires!*'. Our first task, therefore, was to consider the *function* served by decision making in these and other contexts. This led us to think of the function of decision making as part of an attempt to *gain control*, i.e., as an attempt to achieve some desired state of affairs.

Second, the *dynamic character* of the tasks facing decision makers became apparent in the applied contexts that we analyzed. Thus, the decisions could usually only be understood as part of an ongoing process, and the decision problems conformed to Edwards' (1962) classic description of dynamic decision making in that

(1) A series of decisions is required to reach the goal. That is, to achieve and maintain control is a continuous activity requiring many decisions, each of which can only be understood in the context of the other decisions.
(2) The decisions are not independent. That is, later decisions are constrained by earlier decisions, and, in turn, constrain those that come after them.
(3) The state of the decision problem changes, both autonomously and as a consequence of the decision maker's actions.

In an important later paper, Rapoport (1975) provided a more formal definition, but this did not change the general meaning of dynamic decision making compared to Edwards' original definition.

However, the three characteristics mentioned by Edwards did not fully capture the essence of the decision problems we studied. We also had to add another feature:

(4) The decisions have to be made in real time.

Thus, in these decision problems, it was not sufficient to make correct decisions and to make them in the correct order; the decisions also had to be made at the correct time. In short, dynamic decisions are decisions in *context* and in *time*. This means that the decision

maker must consider the consequences of each decision for future decisions, that he or she is constrained by earlier decisions, and that he may also sometimes be able to correct problems caused by earlier decisions in later decisions. The latter feature of dynamic decisions has also been emphasized by Hogarth (1981).

The real-time character of the decisions creates special problems for the decision maker (Brehmer 1990a). First, decision makers are not free to make decisions when they feel ready to do so. Instead, they have to make the decisions when the environment demands decisions from them. This introduces an element of stress into dynamic decision making. To cope with this stress, the decision maker must try to find a strategy that helps him or her not only to control the dynamic task, but also his or her working situation. One way to achieve this is to lower the aspiration level, and try to control the task at some lower, and less stressful level which does not require as many decisions. Thus, dynamic decisions may often be a compromise between a good strategy for controlling the decision task, and a strategy that enables the decision maker to exert some measure of control over the rate at which he has to make decisions. This kind of compromise is observed in process operators where it is often found that they do not try to run their processes close to the limits that would yield the highest output simply because this would require to much work on their part (Bainbridge 1979). The general idea that cognitive behaviour is a compromise between the demands of the task and the need to conserve one's cognitive resources goes back to Bruner et al. (1956) and their careful analysis of the processing demands introduced by different forms of concept attainment strategies. The work of Bruner et al. must be seen as one of the earliest attempts to study dynamic decision making, although it is generally not thought of as such.

Second, both the system that the decision maker seeks to control and the means that he uses for control must be seen as a *processes*. Indeed, we may define the task facing the decision maker in dynamic decision making *as the problem of finding a way to use one process to control another process* (Brehmer and Allard 1991a).

A third consequence is the need to consider the various time scales of the dynamic decision task and to make sure that all the relevant time scales are monitored (see Brehmer, 1991b, for a general discussion of time scales in dynamic tasks, and A. Brehmer, 1989, for an example from intensive care).

All this may seem very abstract, but it is, perhaps, easier to understand if we consider one of the experimental tasks that we have developed for the study of dynamic decision making (Brehmer and Allard 1991a). This task requires the subject to assume the role of a fire chief, charged with the task of fighting forest fires. Even though we have large areas of forest in Sweden, we cannot create a real forest fire every time we want to do an experiment. We have therefore developed a computer simulation of forest fire fighting, and the subject thus fights simulated forest fires. In the simulation, the subject receives information about fires from a spotter plane via an information system. He then sends out commands to his fire fighting units (FFUs) via a command system by typing instructions on the keyboard. The FFUs report back to the fire chief (the subject) about their actions and location, i.e., whether they are en route to some destination, fighting a fire, or are simply inactive and waiting for new commands. On the basis of these reports and further information from the spotter plane, the subject then issues new commands until the fires have been extinguished. The progress of the simulation is controlled by the computer's clock, which does not stop and wait while the subject thinks or makes decisions.

This description demonstrates that the problem facing the subject in the simulation has exactly the characteristics of dynamic decision making described above. Thus, (1) it requires a series of decisions; (2) the decisions are not independent; as resources are committed by one decision, they are unavailable in later decisions, and current decisions are similarly constrained by earlier decisions, even though it is also possible to correct errors in earlier decisions; (3) the state of the problem changes, both autonomously (fires start and burn according to prevailing meteorological conditions) and as a function of the decision maker's actions (the subject's attempts to extinguish the fire will change its development); and (4) the decisions have to be made in real time (or at least in simulated real time).

Many of the features of real fire fighting are omitted from this simulation. This is done intentionally to simplify matters for the subject, so as to cut down the time required for the experiments. It is presumably of little consequence for the results of our experiments, for they are not designed to study fire fighting as such (had this been our purpose, we would not have stayed in the laboratory). Instead, we are interested in the more general problem of how subjects cope with

spatio-temporal processes of the kind exemplified by the fire fighting task.

It is not difficult to think of a host of decision problems that fit the general description above. We have already mentioned military contexts, emergency management, patient care in hospitals and modern process plants, but such problems also arise in company management, and in our personal decision making in everyday life. Despite this, dynamic decision problems have not received much attention from psychologists. There are at least two reasons for this. The first is that it is difficult to cast these problems into the standard normative theory of decision making that has guided so much psychological research in this field, and to find analytical solutions to these decision problems (Rapoport 1975; Slovic et al. 1977), and the second is that they require new experimental methods.

The problem for standard decision theory lies in the decision trees for these decision problems which rapidly become impossibly large and, despite this, such decision trees still fail to incorporate real time.

As for methods, experiments on dynamic decision making simply cannot be made with ordinary paper and pencil methods because of the dynamic and interactive nature of the experimental tasks. Experimentation in these areas is impossible without interactive computer programs, and computers powerful enough for this kind of work have only recently become standard equipment in psychological laboratories.

Despite these problems, there have been some attempts to study dynamic decision making. We now turn to a brief review of this work.

Earlier approaches to the study of dynamic decision making

Dynamic decision making was introduced to psychologists by Edwards and by Toda in 1962. As noted by Rapoport (1975), the main purpose of Edwards' paper was to extend SEU theory from static problems to dynamic problems, and to propose a general research strategy, the normative-descriptive approach, for its study. This strategy involves constructing an 'ideal decision maker' similar to the 'ideal observer' in psychophysics and then comparing the behaviour of real decision makers with that of the ideal decision maker. If systematic discrepancies are found, these may be interpreted as evidence of the

limitations imposed on the actual decision maker, such as limitations of memory and limitations in the ability to predict the effects of one's decisions into the future. Hence it would become possible to model the human decision maker as a case of constrained optimization. Rapoport (1975) provides examples of how this can be done for some sequential decision problems.

A major problem with this approach is that it is impossible to find analytical solutions for most dynamic decision problems (Rapoport 1975). An added problem, also discussed by Rapoport (1975), is the 'flat maximum problem'. That is, for many dynamic problems, even wide discrepancies from the optimal strategy have very little effect on the outcomes of the decisions. Hence, it is difficult to find the limitations imposed on the decision maker using this strategy. It is therefore not surprising that the normative-descriptive strategy has not attracted the wide following among those who study dynamic decision making that it has among those who focus on static problems.

Toda's (1962) paper introduced a one-person game, 'the fungus eater', as a paradigm for the study of dynamic decision making. However, it proved possible to find analytical solutions only for very simplified versions of the game, and its empirical potential was never fully realized. The general approach of constructing a computer simulation as a device for studying dynamic decision making was prophetic, however.

Subsequent analyses of dynamic decision making have largely abandoned normative ambitions and concentrated on empirical work. Before turning to a review of these results, I shall discuss a general framework for the study of dynamic decision making.

Towards a framework for studying dynamic decision making

As noted above, it is not feasible to analyze dynamic decision problems in terms of classical decision theory. We therefore need an alternative framework to guide research. Here, the observation that decisions are made to achieve some overall goal is helpful. As noted above, a general formulation of this observation is to think of the goal of decision making as that of *achieving control*: that is, that decisions are made to achieve some desired state of affairs, or to keep a system in some desired state.

This is not an original observation. Broadbent et al. (1986) made it earlier in their work with dynamic problems, as did Rapoport (1975) and Mackinnon and Wearing (1985). It suggests a general framework for research on dynamic decision making: that of *control theory*. Control theory is a branch of engineering, and this seems quite appropriate; from the functionalist psychological perspective adopted by the present author, it makes more sense to think of man as an intuitive engineer than as an intuitive scientist.

This is, of course, not the first attempt to use control theory as a theory of behaviour. The theory has been quite fashionable in studies of manual control (see, e.g., Kelley 1968; Powers 1971). However, like Bainbridge (1981), we have not found the mathematical framework of control theory particularly helpful. There is no need to repeat Bainbridge's excellent discussion here, except to mention that a major problem with control theory as a psychological framework lies in its inability to handle the kind of input descriptions that are natural in the analysis of human perception: patterns and Gestalts. Instead, it requires that the input be described in terms of measured individual signals.

Therefore, control theory is useful mainly as metaphor. Its usefulness stems from the fact that it specifies the general conditions for control of any system, regardless of whether control is exercised manually or by some automatic controller. There are four general preconditions:

- there must be a goal (the goal condition),
- it must be possible to ascertain the state of the system (the observability condition),
- it must be possible to affect the state of the system (the action condition),
- there must be a model of the system (the model condition). [1]

[1] It should be noted that this is a purely normative statement, and it applies to people as well as machines. However, in case of machines it seems inappropriate to say that the machine has a model of the system it controls. It may be better to say that the machine is a model of the system it controls following the expression used by Conant and Ashby (1970) in the title of their classic paper: 'Every good regulator of a system must be a model of that system'. From a connectionistic perspective, the same expression may apply to humans. Both in the case of human and automatic controllers, the model can, of course, take may forms, but all of them will have the effect that the controller will behave as if he/she/it has a model with some specified relation to the system it controls, and the level of control will depend on the extent to which the model includes the important aspects of the system to be controlled

Properly speaking, two of the conditions pertain to the system to be controlled, whereas the other two are properties of the decision maker. Thus, observability and action possibilities are properties of the system, whereas the goals and the models are properties of the decision maker. The most general formulation of the problem for research, therefore, is that it is concerned with people's formulation of goals and models as a function of the observability and action possibilities of the system to be controlled. This is not quite so straightforward as it seems, however. First, the concepts of observability and action possibilities are abstract, and it is a real research problem to ascertain what constitutes observability and action possibilities for a human decision maker. Second, there is also the problem of *strategy*. A system may be controlled by means of either a feedback strategy or a feedforward strategy (or some combination). In feedback control, the controller uses only the current information about the system, and takes this as information about the actual state of the system. In feedforward control, the controller uses a model of the system to predict its state and to select the appropriate control inputs. Feedback control is possible when there are no significant feedback delays in the system, and it is useful when the system to be controlled changes over time, so that no stable model of the system can be constructed. Feedforward control requires that the system is stable, so that the model remains valid, but it can then use a model to overcome the problems of feedback delays. It is, of course, also possible to use a combination of these two strategies. Modern applications of control theory in the automatic control of systems often rely on a model to produce the actual control inputs, while the feedback information is used to update the model. Brehmer (1990) has hypothesized that feedback control, which is cognitively simpler, would be the preferred mode of control in dynamic decision tasks. The results presented in Brehmer (1990) support the hypotheses, but as we shall see, this may be an oversimplification.

The strategy issue is important because the kind of model that is needed for control depends on the strategy adopted. Thus, one cannot study the subjects' models without first assessing their strategy.

A general framework for research is only a first step. We also need an experimental paradigm that allows us to study dynamic decision making empirically. This paradigm should allow us to define observability and action possibilities in psychologically meaningful ways, and

to observe or infer the subjects' goals and mental models. We now turn to this problem.

Computer-simulated microworlds: A research paradigm for the study of dynamic decision making

Dynamic decision making can be studied in the field, but the complexity of most dynamic systems outside the laboratory prohibits such an approach. For example, to study the mental models of process operators, the researcher might spend many months trying to learn the process of which the operator is assumed to have a model. Moreover, even if one understands the process, the task of collecting data is formidable, and some form of experimental intervention is usually required to obtain useful data (see Hoc, 1989, for a good example). It is therefore not surprising that no theory of dynamic decision making has emerged from field studies so far. It is therefore necessary to complement field studies with more laboratory-based methods. The remainder of this paper will be concerned with some of these methods and the results obtained with them.

As noted above, dynamic decision problems are characterized by the fact that the state of the problem changes both autonomously and as a consequence of the decision maker's actions, and that the decisions have to be made in real time. To capture these features in an ordinary paper and pencil experiment is not feasible, and laboratory research on dynamic decision making therefore had to await the arrival of a new research tool: the *laboratory computer*.

Laboratory computers make it possible to create dynamic simulations and to study how subjects interact with such simulations. We have termed such simulations 'microworlds' (Brehmer and Dörner, in press). Such microworlds simulate some of the essential features, but not all the detail of some dynamic system; they present a 'wood cut' (*Holzschnitt*), to use Dörner's (1991) term, of systems such as a small town, a developing country, an industrial process or a forest fire and charge the subject with the task of controlling this system for some period of time.

Such microworlds are designed to reflect three intuitively striking characteristics of real-world decision problems: *complexity*, *dynamics*

and *opaqueness*. They are complex in the sense that they require the subjects to consider many different elements, e.g., many different. – and possibly conflicting – goals. They are dynamic in some or all of the four senses described above. They are opaque in that they do not reveal all their characteristics automatically to the subject, thus requiring him to form and test hypotheses about their state and characteristics.

Microworlds differ in the extent to which they emphasize one or another of these characteristics. Microworlds such as Lohhausen (Dörner et al. 1983) or Moro (Dörner et al. 1986) developed by Dörner and his associates, who work in the German tradition of research on '*komplexes Problemlösen*', are characterized by high complexity and opaqueness, whereas microworlds such as DESSY (Brehmer and Allard 1991a) and NEWFIRE (Lövborg and Brehmer 1991), which are designed to study decision making, play down complexity and opaqueness and emphasize dynamics. One practical consequence of this is that the progress of the latter kinds of simulations is controlled by the computer's clock. In simulations of the former kind, the subjects make a number of moves, after which the simulation stops and waits for input from the subject before going on to the next state.

A recent experiment by Brehmer et al. (1991) suggests that this difference in simulations has the same effects as a variation in rate of change in the process to be controlled: overall performance is better in a discrete time simulation, but the relative effects of a factor such as feedback delays remain the same. This suggests that the results from the two kinds of simulation are compatible, and I shall ignore the distinction between discrete-time and continuous-time simulations in the remainder of this paper.

Experiments with microworlds

As noted by Brehmer and Dörner (in press), the use of computer-simulated microworlds requires us to rethink the way we do experiments. In such microworld experiments, the actual decision problems are not directly controlled by the experimenter. Instead, they depend on the behaviour of the subject. The experimenter can create a system with certain characteristics, i.e., a given level of complexity, but the trajectory through this system is determined by the subject: the state

of the system at any given moment is dependent both on the general characteristics of the system *and* on the decisions made by the subject. Thus, we need to think of microworld experiments in cybernetic terms rather than in the linear causal terms usually preferred by psychologists, a point also made by Rapoport (1975) in an earlier SPUDM major paper on dynamic decision making.

The independent variables of these experiments, then, are *system characteristics*, such as complexity and feedback structure, but not specific aspects of individual decision problems, such as probabilities and utilities. Therefore, it makes little sense to study isolated individual responses. Instead, we must focus on the decision maker's *tactics* and *strategies*. It is only at this level that we can hope to find regularities in the data from our experiments. This poses special problems of data reduction. It may well be necessary to use simulation to identify these strategies: some actual trajectory is compared to those expected on the basis of some ideal strategies, and the strategy that exhibits the best fit is taken as that which actually describes what the subject did (see Reichert and Dörner, 1987, for an example).

Data reduction is indeed a very severe problem in these experiments because an experimental run will produce thousands of individual observations. Moreover, the trajectory of each individual in these simulations may be unique. This has led Dörner (1991) to take the view that we approach these simulations as case studies. That is, we must study the individual subjects in great detail, describe their individual trajectories and use these to infer the psychological processes. To do so, we must construct computer simulations of the subjects' behaviour. This is reminiscent of Newell and Simon's (1972) approach to the study of problem solving. So far, however, there are no examples of this strategy in the study of dynamic decision making (but see Wearing et al. 1991).

Despite the individual trajectories on the detailed level, it is often possible to obtain stable data at a higher level of aggregation, provided that some care is taken in the data reduction process. Specifically, it is necessary to find the right level of analysis, i.e., a level that reflects the subjects' strategies. This is illustrated in the highly stable and psychologically meaningful relations among various performance measures and measures showing how subjects used their resources in the fire-fighting simulation (e.g., Brehmer 1990).

Research problems in the study of dynamic decision making

Research on dynamic decision making has followed two different approaches: an individual differences approach and a more standard experimental approach.

The individual differences approach has been pursued mainly by Dörner and his associates. In the typical experiment, all subjects perform in the same complex simulation and are then divided into two (or more) groups according to how they performed in the simulation. These groups are then compared with respect to their performance on some test, or with respect to their behaviour in the simulation, to find possible explanations for the differences in performance (see Dörner et al., 1983, and Stäudel, 1987, for examples).

In the experimental approach, subjects' performance in simulations with different characteristics, e.g., feedback delays, side effects or complexity, is compared (see Brehmer and Allard, 1991a, Funke, 1990, for examples).

Such experiments are undertaken not only to investigate the effects of system characteristics, but for a variety of purposes. For example, Broadbent and his associates have used different dynamic simulations to study the issue of awareness and to formulate alternatives to the serial, STM–LTM-model of human information processing (e.g., Berry and Broadbent 1987, 1988). However, I shall not discuss these uses because they have little to say about dynamic decision making.

In the remainder of this paper, I shall review some of the empirical results obtained with complex dynamic microworlds.

Empirical examples

This section is in two parts. The first reviews the results from the individual differences approach and the second those from the experimental approach.

The individual differences approach

As already noted above, the principal method in this approach has been to have a group of subjects perform in the same microworld, and then divide them up into two extreme groups, one that succeeds and

one that does not. These groups are then compared with respect to their behaviour, and/or their scores on various psychological tests.

Research following this approach could have two different goals: to *find ways of predicting behaviour* in dynamic decision tasks, so that it would then be possible to select good decision makers, and to *identify the demands* made by these tasks by comparing the behaviour of the subjects who are successful in controlling the microworld with that of those who are not so successful (a very useful introduction in English to this work is provided in two papers by Funke, 1988, 1991).

Prediction of performance

Results from various simulations show that there are wide individual differences in performance. Despite this, it has proved difficult to find any significant correlations between performance in microworld experiments and test performance. Performance in the microworld experiments does not seem to correlate with any standard psychological test, be they tests of intelligence or tests of personality (Dörner et al. 1983; Stäudel 1987). There is some evidence that performance in microworlds may correlate with intelligence, but only when the opaqueness is removed (Putz-Osterloh and Lüer 1981; Schoppek 1991). Some results indicate that a new variable defined by Dörner and his associates, and operationalized by Stäudel in a questionnaire, might correlate with performance. This variable was called 'heuristic competence', and may be defined as a general competence for coping with complex systems.

Stäudel (1987) found that subjects high in heuristic competence performed better in Moro than subjects with low heuristic competence. Moro is a complex system designed by Dörner (Dörner et al. 1986) which requires the subject to assume the role of an adviser to a tribe in southern Sahara, and advise them on how to improve their conditions during a simulated 20- or 30-year period.

Schaub and Strohschneider (1989) obtained further evidence for usefulness of Stäudel's concept in a study in which they compared the behaviour of one group assumed to have high heuristic competence, managers from large German and Swiss corporations, with a group assumed to have low heuristic competence, psychology students. Although the differences in performance were not dramatic, managers clearly performed better in the Moro simulation. These differences were not due to any greater knowledge about the domain covered by

Moro, i.e., the economics and ecology of developing countries (what Dörner has called 'epistemic competence'). Instead, the managers performed better because they approached the decision task in a different way. The managers were less activistic than the students, i.e., they made fewer decisions, they collected more information before they made their decisions, and they checked on the results of their decisions before making new decisions. This is exactly what is required to perform well in a complex dynamic system (the behaviours reflect what Dörner, 1989a, has called 'grandmother rules') and may be seen as an expression of greater heuristic competence on the part of the managers, i.e., they had learned some general ways of coping with the demands made by complex dynamic systems.

To be successful, this approach requires that reliable performance measures can be found. A study by Strohschneider (1986) suggests that this is problematic, at least for Moro, one of the commonly used simulations. He found no significant correlations between performance in two consecutive runs with Moro. This suggests that one possible explanation for the failure to find stable correlations between performance in microworlds and test performance may be that the criterion variable lacks reliability. This is hardly surprising. The subjects' performance in a microworld is dependent not only upon their performance but also on the characteristics of the microworld. In some microworld scenarios, a small and seemingly unimportant mistake may lead to serious problems later on, from which the subject cannot recover, even if he follows a good strategy. Thus we should, perhaps, not be too surprised that it has proved so difficult to find correlations between microworld performance and test performance.

The demands that complex dynamic systems make

The second goal of the individual differences approach is to ascertain the demands made by complex dynamic microworlds. The first method here has been to compare the behaviour of extreme groups, e.g., the best and the worst third of the subjects, on psychological tests. The rationale behind this approach is that if we find, for example, that subjects who score high on spatial ability perform better than those who score low on these tests, we might conclude that the simulation in question requires spatial ability. As noted above, the search for such correlations has not been very successful.

The second method has been to analyze the differences in subjects'

behaviour in the simulation for extreme groups. This has not produced any surprising results. It seems that subjects who collect more information, who collect it more systematically, who construct adequate goals, who evaluate the effects of their decisions, and who generally behave in a systematic fashion tend to perform better than those who do not (Dörner et al. 1983). Thus we see the 'grandmother rules' in operation once more. However, the results of Dörner et al. (1983) point to one important conclusion: those subjects who behave in a way that makes it more likely that they will acquire a good model of the task also learn to control the task better (see also Funke and Müller 1988).

It has not yet proved possible to improve performance of subjects by training them in the general methods for handling complex systems. For example, in the Lohhausen study, Dörner et al. (1983) compared two different training programs, one designed to give general strategic insight and one designed to teach useful tactics for coping with complex dynamic systems. Subjects in the control group received no special instructions. There were no differences in the performance of the three groups. Because the training programs can be seen as an attempt to induce heuristic competence, these results may, at first, seem to be in conflict with those of Stäudel and of Schaub and Strohschneider with respect to heuristic competence. However, it may well be that the skills involved here cannot be taught in the abstract, i.e., as a set of general rules, but that they can only develop in conjunction with actual practice in controlling systems. Such practice may, for example, teach the decision makers *when* and *how* to apply a given rule. A simpler, and more trivial explanation, for the results of Dörner et al. is, of course, that the training methods were not appropriate, and before drawing any more definite conclusions, we will need to investigate the effects of a wider range of training programs. However, the issue of training is important, and deserves further study.

Simulations such as Moro and Lohhausen are designed to mirror at least some of the characteristics of real systems such as developing countries (Moro) and towns (Lohhausen); it, therefore, seems reasonable to expect that *epistemic competence*, i.e., domain knowledge, would be a powerful factor in these experiments. This aspect has, however, not received much attention. Hasselrot and Brehmer (1991) compared a group of university students taking courses in the agricul-

ture, ecology and economics of developing countries with a group of psychology students. They found initial differences favouring the students with high epistemic competence, but in the long run both groups produced similar dismal conditions in Moro, because the subjects with high epistemic competence did not exhibit any evidence of heuristic competence. Specifically, they did not check on the results of their decisions as they should have done, and they thus ignored the development of one of the crucial variables in Moro: the ground water level. These results suggest, then, that both epistemic and heuristic competence are needed, but if a subject were to have only one of these, heuristic competence would be more useful than epistemic competence.

The pathologies of decision making

Although the study of successful subjects seems to have uncovered no more than that these subjects, for whatever reason, exhibit the kind of behaviours that would be expected to be useful in these kinds of systems, the study of unsuccessful subjects has led to some novel and interesting results. These studies have uncovered a number of general maladaptive behaviours that could be called *pathologies of decision making*. These pathologies are not necessarily the cause of failure, but are observed as subjects try to cope with their failures (Dörner 1980). Dörner (1989a) theorizes that these behaviours may be the consequence of feelings of low competence that result from failure. They then lead to even greater failure in a self-reinforcing cycle.

The first of these pathologies is called 'thematic vagabonding', and it refers to a tendency to shift goals. Thus, a subject in, say, Moro, may start out dealing with the water supply. When this fails to solve all the problems, he may then concentrate on improving education, then shift back to working on the water supply, and then move on to do something about storage, and so on. Thus, subjects who exhibit thematic vagabonding shift from one subgoal to another but fail to work on the problem as a whole, which is what is needed in a complex system such as Moro.

The second pathology is called 'encystment'. It involves sticking to one specific goal that the subject feels competent to achieve. Thus, a subject who has great faith in education perseveres with this method as the method which is to solve all problems in Moro (see Dörner et al., 1983, for examples).

A third pathology is a general refusal to make any decisions at all.

A fourth pathology involves blaming others for one's failures. (Many of these programs have simulated others who have to carry out the subject's commands.) One of the most likely candidates here is, of course, the simulation itself, a response similar to that of blaming 'the system' for one's failures.

A fifth pathology is that of delegating responsibility that cannot, or should not, be delegated.

There is, of course, also the pathology of not delegating responsibility which should be delegated. For example, in systems with feedback delays, it is often useful to delegate responsibility from a central level (which is subject to great delays) to more local decision makers (who are not subject to as severe delays in information). This is an important principle in military organizations (see Brehmer 1988). However, Brehmer and Allard (1991b) found that subjects in the fire-fighting simulation did not use delegation as a means to combat the effects of delays.

It is interesting to note that these pathologies can be divided into two groups. The first group comprises the first two pathologies. These pathologies can be seen as failures of the goal formulation process. The second group comprises the last three pathologies, which may be seen as signs of a refusal to learn from experience. Thus, the first of the three pathologies in this group, refusal to make decisions, clearly means that the subject will not receive any new experience to learn from, while the last two involve rejecting whatever experience could have been used for learning.

The results with respect to pathologies so far consist mainly of a set of examples. We do not know the relative frequency of these pathologies, nor do we know if there are individual differences among people with respect to which of these pathologies they will exhibit. Indeed, we do not even know if the pathologies are common. This is an important problem for future studies.

The experimental approach

Research following the experimental approach aims to assess the effects of the characteristics of the system on the behaviour of subjects in dynamic tasks in order to enable the investigator to make inferences about how the subjects develop mental models and formulate

goals. This requires a taxonomy of the task that would describe it in terms of those aspects that affect performance. The problem of taxonomies for complex dynamic microworlds has been discussed by Brehmer (1990c, 1991a) and Funke (1990). However, I shall not dwell on the general problems of taxonomy here but use a simple set of concepts proposed by Brehmer and Allard (1991a) to organize the results obtained so far.

The characteristics of dynamic tasks

A taxonomy of dynamic systems is interesting only if it helps predict the behaviour or performance of the subjects. This means that the taxonomy is a first step towards a theory. It is therefore important that the set of dimensions chosen for the taxonomy has some psychological meaning.

Brehmer and Allard (1991a) have proposed such a set of dimensions to be used to describe dynamic tasks and to guide experimentation. They proposed six dimensions.

The first of these is *complexity*. This is a difficult concept, for it cannot be given an absolute definition, nor does it have an obvious metric. Instead, it is a relative concept, and it must be defined in relation to something or someone for which the system is complex. In the present context, our interest is in control, and it is therefore reasonable to define complexity in relation to the capacity of a controller. This is the approach proposed by Ashby (1956) in introducing the concept of *requisite variety*. In this case, the controllers of interest are humans, and it is in relation to their limitations that a task may be seen as complex. One of the fundamental limitations of humans is the number of items they can process at a given time. It therefore makes sense to define complexity in terms of the number of system elements and their relations. However, there can be no simple relation between the number of elements of a system and difficulties in controlling it. First, not all elements and their relations may be relevant to a given control task. Thus, we must consider the subsystem relevant to a given task when trying to assess the relevant complexity of the system. Second, we must consider the nature of the relations among the elements and the nature of the elements. For example, if a system has many elements and the relations are negative feedback loops, the system will be very stable, and will not develop in a

catastrophic direction regardless of what the decision maker does. If the relations are positive feedback loops, on the other hand, even a system with few elements may be hard to control. A mixture of positive and negative feedback loops may lead to system behaviour that is hard to understand and predict, and may make control very difficult (see Mackinnon and Wearing, 1980, for a discussion of these problems). Moreover, not all elements may contribute equally to the difficulty of a system. Brehmer and Allard (1991a) distinguished four kinds of elements: goals, processes to be controlled to reach the goals, control processes, and side effects. This was an attempt to distinguish among different elements in relation to the subjects' purposes. Whether this also makes empirical sense is still an open question.

The second dimension is *feedback quality*. In dynamic tasks, all information is feedback information, but this feedback information may nevertheless vary in quality, leading to variation in the observability of the system.

The third concept is *feedback delays*. This is the interval between the point in time when an action is taken and the time when information about the effects of that action arrives. In complex systems, feedback delays are inevitable. The concept of delay is a complex one, however. First, feedback delays may occur anywhere in the feedback loop. Thus, they may occur in the response to commands (what control engineers call *dead time*), in the execution of commands (usually called the *time constants* in control theory), or in the transmission of information about the results of the action taken (*feedback delay*). Second, the information about these delays may vary. Thus, the concept of observability is relevant here, too.

The fourth concept is *rate of change*. This refers to the rate of change in the process to be controlled. Dynamic tasks vary widely in their rate of change, from the very rapid tasks facing a fighter pilot making a low level attack, to the very slow tasks facing the minister of finance of a country trying to control the economy of that country. Presumably, there is a range of values here with which a person may feel comfortable. This is probably the range within which it is possible to control the tasks on the basis of direct feedback. As tasks become very slow, the dynamic character may well be ignored, and the task seen as a series of independent decisions, or, at best as a set of sequential decisions. As the rate of change becomes very fast, it is no longer possible to control the system on the basis of feedback. Instead,

the controller must develop some means of feedforward or open loop control.

The fifth concept is *the relation between the characteristics of the processes to be controlled and those of the process used for control*. This defines the strategy for the task. For example, fire fighting of the kind represented above is distinguished by the fact that one linear process, the fire-fighting process, is used to fight another linear process, the spreading fire, but these processes differ in their rate of change: the fire grows faster than the fire-fighting units are able to extinguish it. Therefore, to succeed in extinguishing the fire, the subject has to attack it massively and rapidly, i.e., the subject must send so many fire-fighting units to a given fire that they cover at least the area that will be on fire when the units arrive at the fire. Matching the number of fire-fighting units to the fire as it is when the units are sent out will not be sufficient. The subjects quickly learn the elements of this strategy in the fire-fighting simulation (see, e.g., Brehmer and Allard 1991a). This concept encompasses both the '*Eigendynamik*' of the system, i.e., the extent to which it changes autonomously, and the feedback structure, i.e., whether the feedback is positive or negative. Both of these will be aspects of the characteristics of the process to be controlled. Thus, this concept will need to be differentiated into a number of subclasses.

The sixth concept refers to *the extent to which the decision making power in the system can be delegated*, or distributed, among the persons in the system. This is especially important in systems where there are long delays. In this case, delegating the decision making to lower levels in the system that are closer to the actual events may increase efficiency, a fact very well recognized in military contexts (see Brehmer 1988).

This list of concepts is, of course, only a preliminary one, to be replaced by a new list as research progresses. From a systems point of view, it is a problem that the concepts are not independent. For example, there is a close relation between the relation between the characteristics of the process used for control and the characteristics of the process to be controlled and the concept of feedback delay. Our problem, however, is not only to characterize systems, but to find psychologically meaningful concepts for describing systems. Whether the one concept can be subsumed under the other is an important question for future research.

Experimental results

Research with computer-simulated microworlds is still in its infancy, and there are as yet few well-established findings concerning the effects of various system characteristics. However, even the small body of results that are available shed some interesting light on people's ability to control system. In this section, I shall review these results, using the list of concepts discussed above to organize them.

Complexity

Despite the central position of the concept of complexity in the German tradition of '*Komplexes Problemlösen*', there have been few attempts to study the effects of complexity systematically. Instead, the strategy has been to create a simulation that can be said to be complex, and then to use this to compare the behaviour of different groups of subjects as described above.

One reason for the lack studies of the effects of complexity may well be that there is something trivial about this factor as an independent variable in experimental research. Because there is no metric for complexity, it must be defined in relation to the capacity of the subjects. That is, either the subjects can mobilize the requisite variety needed to control the task or they cannot. Moreover, their ability to produce requisite variety cannot be seen as a fixed property. Instead, it is likely to develop with practice. However, even though demonstrating the effects of complexity may be uninteresting, assessing the variety a person can produce, to track its development, to find ways of speeding up this development or to provide means for augmenting it, are far from trivial research problems.

Because there is so little experimental work on the effects of complexity factors, it is difficult to interpret some of the results that we have. For example, there is some consensus that people have problems taking the side effects of their decisions into account, a conclusion based on analyses following both the individual differences approach (e.g., Dörner et al. 1983) and the experimental approach (Müller et al. 1988). This is also confirmed by Sterman (1989), who showed that subjects had problems understanding the feedback between their responses and the environment in an economic simulation. Specifically, his results show that the subjects did not understand that their demand for capital would affect the total demand for capital, leading to at least temporary problems in getting the capital

they needed. This may, of course, be seen as a side effect, but it is also evidence that one of the fundamental characteristics of a dynamic decision task: that the state of the task changes both autonomously and as a consequence of one's own behaviour, is difficult to grasp.

It is not clear whether the problems people have with side effects should be seen as an inherent characteristic or whether the results simply express the difficulties naive subjects have when working with systems about which they know little. That is, the subjects in the microworld experiments may simply not have had enough experience to learn about the important side effects of the decisions they have to make. This raises the problem of the extent to which decision makers can be trained to consider side effects to a greater extent, and what kinds of side effects they may learn to consider in their decisions. A limiting factor here is, of course, the decision maker's attentional resources. Because these are limited, we would expect at least some side effects will be ignored in most complex situations.

Another problem making it difficult to assess the effects of complexity arises from the interaction between different factors in the system. This is illustrated by some results of Mackinnon and Wearing (1980). They failed to demonstrate any effects of complexity defined in terms of the number of elements in the system when the system was characterized by strong negative feedback loops. Under these conditions, the system becomes inert and self-stabilizing. That is, even though the strategy used to control the system is inappropriate, the negative feedback loops of the system may save the decision makers from the worst consequences of their inappropriate strategies. On the other hand, they will also be unable to achieve the kind of change that they strive for.

The results from Dörner's studies suggest that subjects have considerable problems finding adequate goals in complex simulations (Dörner et al. 1983). However, it is hard to decide whether this is due to the complexity of the simulation. There are no direct studies on, for example, the effects of the number of goals on subject performance. Again, this is due to the research strategy of constructing a system that is complex enough, and then using it without any experimental variation of the number of goals or the number of goal conflicts. Moreover, in these experiments subjects usually receive only very general and diffuse goals in the instructions (such as to 'care for the welfare of the moros'), and they have to derive a series of more

concrete goals before they are able to act. This emphasizes the close relation between the goal formulation process and the mental models that subjects may have. If these models are inadequate, the goals formulated on the basis of them will be inadequate also. Consequently, it is not possible to decide whether the goal formulation problems observed in, for example, the Lohhausen study are due to inadequate models or to problems in the goal formulation process as such.

Feedback quality

As noted above, microworlds are usually designed to be opaque: they are designed so that the feedback quality is less than perfect. That is, the subjects are not informed about all of the consequences of their actions automatically. Instead, they must often use more indirect information to infer the state of the system, or at least actively search for information about some of the relevant aspects of the system. This requires that they know, and remember, what these aspects are.

There are no experimental results on the effects of feedback quality as such. However, the importance of this factor is illustrated in studies using Moro. Here, one of the main problems is the ground water level, a variable that is not displayed automatically to the subject. To control this factor, the subject must remember to ask for information about it. The failure to ascertain the ground water level is one of the main reasons why so many subjects fail in Moro (see e.g., Hasselrot and Brehmer 1991).

The subjects' ability to develop the appropriate strategy of rapid and massive response to a fire in the fire fighting simulation and their inability to cope with delays in reports from the fire fighting units (Brehmer and Allard 1991a,b) may also be seen as a matter of differences in feedback quality. Thus, the delays caused by the time constants of the system, which are due to the speed at which the fire fighting units move, are directly observable, but the report delays have to be inferred.

Feedback delays

Feedback delays are interesting for at least two different reasons. First, they are inevitable in most complex systems, for everything takes time, and a decision maker cannot expect to have the results of his actions displayed immediately. Whether he can cope with these delays

is a central feature of his ability to cope with complex dynamic systems. Second, the effects of feedback delays give us information about the subjects' general strategies. For example, if they do not compensate for the delays, this suggests that they employ a feedback control strategy, but if they compensate adequately for the delays, this suggests some form of feedforward strategy, or at least some open loop strategy.

Feedback delays may occur in different parts of the feedback loop. Depending on where in the feedback loop they occur, different models are required to cope with the delays. Thus, dead time and time constants will require a model that makes it possible to predict when an action is needed before the need has arisen. Delays in the reports about the effects, on the other hand, require a model which allows prediction of when a given action has taken effect and when the resources will be available again. Developing the former kind of model seems inherently more difficult, unless there is some regularity in the need for decisions about actions. We would therefore expect that dead time would cause more difficulties for decision makers than report delays do.

A number of studies show that feedback delays have negative effects on performance in a variety of complex dynamic systems (Brehmer and Allard 1991a,b; Dörner 1989b; Reichert and Dörner 1987). However, as noted above, the negative effects may be limited to delays that have to be inferred (Brehmer 1990).

As expected, there are also differences in the effects of delays that occur in different locations in the feedback loop. In the fire-fighting simulation, dead time causes greater problems than report delays (Brehmer and Allard 1991c). This is hardly surprising. If the subjects fail to compensate for both forms of delays, as they seem to do, they send fewer fire-fighting units to the fire under dead-time conditions (because the fire will spread even further during the dead time) than they will under report-delay conditions, where the subjects have an adequate picture of the fire at the moment of decision. The predicted differences in use of resources between the two kinds of delay conditions were also found by Brehmer and Allard (1991c). Similar results were obtained when the subjects were informed about the delays.

As noted above, we would expect that dead time would be more difficult than report delays in general (this is true also in engineering

applications of control theory to real systems). Once the need for a decision is apparent, on the other hand, dead time will have very much the same effects as report delays, provided that the process to be controlled does not change too rapidly in relation to the dead time. That is, in both cases, the main effect will be that it will take some time before the decision maker can see the effects of his actions.

Subjects do not use the strategy of delegating decision-making power to local units to help them cope with the delays. Thus in experiments using the fire-fighting simulation, Brehmer and Allard (1991a,b) gave subjects the opportunity to delegate decision making to the fire-fighting units. Specifically, they could delegate the decision when to start fighting a fire to the (simulated) fire-units commanders, so that they would not have to wait for a report from them about their location to decide whether they needed to change their destination to where the fire was most threatening. However, subjects did not delegate decision-making power more when there were delays. Indeed, they were generally unwilling to delegate any decision making to the unit commanders, and they generally became less willing to delegate as the experiment progressed, even though they were not doing particularly well when using their centralized decision-making policy.

The results reviewed in this section suggest that subjects will adopt a feedback-control strategy when the delays will have to be inferred and an open-loop strategy when the delays can be seen to happen (Brehmer 1990). Thus, subjects seem to develop different models of the task depending on the observability. The results also show that subjects seem to adopt the same feedback control strategy for dead time and report delays when they are told about these delays as when they are not told about them and when sufficient information to infer the delays and their causes is provided in displays in the experiment. This shows that observability from the subjects' point of view is not the same as from a control engineer's point of view.

Rate of change

The effect of increasing the rate of change in a dynamic decision task is to increase the work load of the subject. Predictably, this leads to worsening performance (Brehmer and Allard 1991d; Brehmer et al. 1991).

The relation between the characteristics of the process to be controlled and the characteristics of the control process

There are no systematic studies of this factor as yet. In most simulations, the control process tends to be a linear one. Perhaps this is typical of man-made control processes. However, we could easily think of exponential control processes, e.g., when one biological process is used to control another biological process. Dörner (1980) reports that exponential processes cause difficulties for the subjects, but this conclusion is not based on any systematic comparison between different processes to be controlled or between different relations between the control process and the process to be controlled. Funke (1990) reports that processes with '*Eigendynamik*' are difficult to control, and more difficult the more pronounced the '*Eigendynamik*'. Again, the difficulty of controlling a process with '*Eigendynamik*' may be due to the characteristics of the control process. Thus, we should reserve judgment about the relative difficulty of different processes until we have investigated its difficulty relative to different control processes.

The fire-fighting simulation of Brehmer and Allard (1991a) uses a linear process (fire fighting) to control another linear process (the fire), where the slope of the latter exceeds that of the former. As shown by Brehmer and Allard's results, subjects master this relation, at least if they are able to see the processes, as they can in the fire-fighting simulation where both the spread of the fire and the movements of the fire-fighting units are displayed to them.

These results also show the close relation between the characteristics of the process to be controlled, the characteristics of the control process and the concepts of time constants and dead time. The slope of the function describing the process to be controlled is obviously related to the time constant of the process. Dead time introduces a nonlinearity in the function describing the control process. This raises the problem of whether it is meaningful to retain the distinctions among dead time, time constants and the functional characteristics of the processes involved. This is a question that can only be answered by further psychological research. Only if subjects react differently in response to the nonlinearity introduced by dead time and that introduced by other factors, is it psychologically meaningful to make the distinction.

Summary

The results reviewed here probably reflect only one aspect of the subjects' performance: the nature of their mental models. They suggest that these mental models do not include all aspects of a complex dynamic task. Thus, side effects and delays that have to be inferred seem to be problematic, and there is little evidence from the subjects' performance in the microworld experiments that they have incorporated them in their models. This is true for delays even when the subjects claim to have detected them (Brehmer and Allard 1991b), and even when they are told about these delays (Brehmer and Allard 1991b,c).

Both side effects and feedback delays are common features of real-world systems. Our conclusion, therefore, would be that we would expect decision makers working with complex dynamic systems to show a suboptimal level of performance.

Conclusions

This paper had three purposes: to introduce a general paradigm for the study of dynamic decision making: that of control theory; to describe a general methodology for the laboratory study of dynamic decision making: that of computer simulated microworlds; and to review results obtained with this paradigm.

Generally speaking, the research described here can be seen as an extension of the ideas presented by Rapoport (the use of control theory) and Toda (the use of computer simulations as research tools), but it differs from Rapoport's and Toda's original papers in that it has abandoned the normative-descriptive comparisons for empirical analysis and psychological theorizing. This is because the analytical analysis of complex dynamic tasks has proved impossible, a conclusion already drawn by Rapoport (1975).

So far, the total body of research is rather small, and much remains to be done. However, the results so far demonstrate that it is possible to investigate even complex dynamic decision-making phenomena experimentally. Moreover, a number of consistent results have been obtained. These results suggest that people may have problems developing the mental models required by complex, dynamic tasks. Some of the results are corroborated by observations from field studies, such as

studies of operators in process plants, who have also been found to have problems with feedback delays and with side effects (e.g., Crossman 1974). This suggests that the results obtained in the microworlds are not simply laboratory artifacts, but that they may express some generally valid characteristics of people trying to control complex dynamic systems.

The present results are limited, however, in one important respect: there has been little need to invoke the concept of decision making. This is, of course, because there have been very few goal conflicts in the tasks used so far. Therefore, the subjects are not required to use their models for different purposes, and the results can generally be interpreted as an indication of the characteristics of the subjects' mental models.

The next step in research on dynamic decision making, therefore, must be to introduce motivational factors in a more systematic way in the experiments (see Omodei and Wearing, 1991, for an attempt to study motivational factors in a nonexperimental way in dynamic decision tasks. See also Huber, 1990a,b, for studies of the effects of cost factors in sequential decision tasks, a form of task bridging between the traditional static decision tasks and dynamic tasks of the kind considered here). For this purpose, we have introduced a new version of the fire-fighting task called NEWFIRE (Lövborg and Brehmer 1991) which makes it possible to vary cost factors in the fire-fighting task systematically. This should make it possible to assess whether subjects are able to modify their strategies in response to economic considerations and not only in response to system characteristics such as time constants.

This should enable us not only to create a firm connection between dynamic decision making and general cognitive psychology, as is done via the interest in mental models, but also between dynamic decision making and the more traditional decision literature via the new possibilities of studying the effects of cost factors.

References

Ashby, W.R., 1956. An introduction to cybernetics. London: Chapman and Hall.

Bainbridge, L., 1979. Forgotten alternatives in skill and work load. Ergonomics 21, 169–185.

Bainbridge, L., 1981. 'Mathematical equations or processing routines?' In: J. Rasmussen and W.B. Rouse (eds.), Human detection and diagnosis of systems failures. New York: Plenum.

Berry, D.C. and D.E. Broadbent, 1987. The combination of explicit and implicit learning processes in task control. Psychological Research 49, 7–15.

Berry, D.C. and D.E. Broadbent, 1988. Interactive tasks and the implicit–explicit distinction. British Journal of Psychology 79, 251–372.

Brehmer, A., 1989. Brännskadevård som ett hierarkiskt organiserat system. Uppsala University: CMD Report No. 4.

Brehmer, B., 1988. 'Organization for decision making in complex systems'. In: L. Goodstein, H. Andersen and S.E. Olsen (eds.), Tasks, errors and mental models. London: Taylor and Francis.

Brehmer, B., 1990a. 'Strategies in real-time, dynamic decision making'. In: R. Hogarth (ed.), Insights in decision making. Chicago, IL: University of Chicago Press.

Brehmer, B., 1990b. 'Dynamic decision making'. In: D. Frey (ed.), Bericht über den 37. Kongress der Deutschen Gesellschaft für Psychologie, Band 2, Kiel, 1990. Göttingen: Hogrefe.

Brehmer, B., 1990c. 'Towards a taxonomy for microworlds'. In: J. Rasmussen, B. Brehmer, M. de Montmollin and J. Leplat (eds.), Taxonomy for analysis of work domains. Proceedings of the first Mohawc workshop. Roskilde. Risö National Laboratory.

Brehmer, B., 1991a. Human control of complex systems: A view from reserch with microworlds. Paper prepared for the Third European Meeting on Cognitive Science Approaches to Process Control. Cardiff, September 2–5.

Brehmer, B., 1991b. 'Time scales, distributed decision making, and modern information technology'. In: J. Rasmussen, B. Brehmer and J. Leplat (eds.), Distributed decision making: Cognitive models for cooperative work. Chichester: Wiley.

Brehmer, B. and R. Allard, 1991a. 'Real-time dynamic decision making. Effects of task complexity and feedback delays'. In: J. Rasmussen, B. Brehmer and J. Leplat (eds.), Distributed decision making: Cognitive models for cooperative work. Chichester: Wiley.

Brehmer, B. and R. Allard, 1991b. Feedback delays in real-time dynamic decision making. Uppsala Psychological Reports.

Brehmer, B. and R. Allard, 1991c. Dead time and report delays in dynamic decision making. Uppsala Psychological Reports.

Brehmer, B. and R. Allard, 1991d. Rate of change in dynamic decision making. Uppsala University, Department of Psychology, Manuscript.

Brehmer, B., R. Allard and A. Jansson, 1991. Discrete time and continuous time simulations in the study of dynamic decision making. Uppsala Psychological Reports.

Brehmer, B. and D. Dörner, in press. Research with computer simulated microworlds: Escaping the narrow straits of the laboratory as well as the deep blue sea of the field study. Computers in Human Behavior.

Broadbent, D.E., P. Fitzgerald and M.H.P. Broadbent, 1986. Implicit and explicit knowledge in the control of complex systems. British Journal of Psychology 77, 33–50.

Bruner, J.S., J.J. Goodnow and G.A. Austin, 1956. A study of thinking. New York: Wiley.

Conant, R.C. and W.R. Ashby, 1970. Every good regulator of a system must be a model of that system. International Journal of System Science 1, 89–97.

Crossman,, 1974. 'Automation and skill'. In: E. Edwards and F.P. Lees (eds.), The human operator in process control. London: Taylor and Francis.

Dörner, D., 1980. On the problems people have in dealing with complexity. Simulation and Games 11, 87–106.

Dörner, D., 1989a. Die Logik des Mißlingens. Strategisches Denken in komplexen Situationen. Reinbek: Rowohlt.

Dörner, D., 1989b. Managing a simple ecological system. Proceedings of the Second European Meeting on Cognitive Science Approaches to Process Control, Siena, 1989.

Dörner, D., 1991. Über die Philosophie der Verwendung von Mikrowelten oder 'Computerszenarios' in der psychologischen Forschung. Working Paper No. 7. Berlin: Projektgruppe Kognitive Anthropologie. Max-Planck-Gesellschaft.

Dörner, D., R. Kreuzig and T. Stäudel, 1983. Lohhausen: Vom Umgang mit Unbestimmtheit und Komplexität. Bern: Huber.

Dörner, D., T. Stäudel and S. Strohschneider, 1986. Moro-Programmdokumentation. Memorandum 23. Lehrstuhl Psychologie II, University of Bamberg.

Edwards, W., 1962. Dynamic decision theory and probabilistic information processing. Human Factors 4, 59–73.

Funke, J., 1988. Using simulation to study complex problem solving. Simulation and Games 19, 277–303.

Funke, J., 1990. Systemmerkmale als Determinanten des Umgangs mit dynamischen Systemen. Sprache and Kognition 9, 143–154.

Funke, J., 1991. 'Solving complex problems: Exploration and control of complex systems'. In: R.J. Sternberg and P.A. Frensch (eds.), Complex problem solving: Principles and mechanisms. Hillsdale, NJ: Erlbaum.

Funke, J. and H. Müller, 1988. Eingreifen und Prognostizieren als Determinanten von Systemidentifikation. Sprache and Kognition 7, 176–186

Hasselrot, A. and B. Brehmer, 1991. Epistemic competence in dynamic decision making. Uppsala University, Department of Psychology, Unpublished manuscript.

Hoc, J.M., 1989. Strategies in controlling a continuous process with long response latencies: Needs for computer support to diagnosis. International Journal of Man–Machine Studies 30, 47–67.

Hogarth, R.M., 1981. Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. Psychological Bulletin 90, 197–317.

Huber, O., 1990a. 'Cognitive processes in multistage decision making'. In: Gilhooly et al. (eds.), Lines of thinking: Reflections on the psychology of thought. Chichester: Wiley.

Huber, O., 1990b. Perceived control in a multistage decision task. Journal of Behavioral Decision Making 3, 126–136.

Kelley, C.R., 1968. Manual and automatic control. A theory of manual control and its application to manual and automatic control systems. New York: Wiley.

Lövborg, L. and B. Brehmer, 1991. NEWFIRE. Roskilde: Risö National Institute. Report M-2953.

Mackinnon, A.J. and A.J. Wearing, 1980. Complexity and decision making. Behavioral Science 25, 285–396.

Mackinnon, A.J. and A.J. Wearing, 1985. Systems analysis and dynamic decision making. Acta Psychologica 58, 159–172.

Müller, H., J. Funke and B. Rasche, 1988. Wechselseitige Abhängigkeiten: Zum Einfluß von Eigendynamik und Nebenwirkungen auf die Bearbeitung dynamischer Systeme. Berichte aus dem Psychologischen Institut der Universität Bonn, 14, No. 1.

Omodei, M.M. and A.J. Wearing, 1991. Decision making in complex dynamic settings: A theoretical model incorporating motivation, intention, affect and cognitive performance. Paper presented at the 13th Research Conference on Subjective Probability, Utility and Decision making. Fribourg, Switzerland, August 18–34.

Powers, W., 1971. Behavior: The control of perception. New York: Addison-Wesley.

Putz-Osterloh, W. and G. Lüer, 1981. Über die Vorhersagbarkeit komplexer Problemlöseleistungen durch Ergebnisse in einem Intelligenztest. Zeitschrift für Angewandte und Experimentelle Psychologie 28, 309–334.

Rapoport, A., 1975. 'Research paradigms for the study of dynamic decision behavior'. In: D. Wendt and C. Vlek (eds.), Utility, probability and human decision making. Dordrecht: Reidel.

Reichert, U. and D. Dörner, 1987. Heurismen beim Umgang mit einem 'einfachen' dynamischen System. Sprache and Kognition 7, 12–34.

Schaub, H. and S. Strohschneider, 1989. Die Rolle heuristischen Wissens beim Umgang mit einem komplexen Problem oder: Können Manager wirklich besser managen? Memorandum No. 71, Lehrstuhl Psychologie II, University of Bamberg.

Schoppek, W., 1991. Spiel und Wirklichkeit – Reliabilität und Validität von Verhaltensmustern in komplexen Situationen. Sprache and Kognition 10, 15–37.

Slovic, P., B. Fishhoff and S. Lichtenstein, 1977. Behavioral decision theory. Annual Review of Psychology 28, 1–39.

Sterman, J.D., 1989. Misperception of feedback in dynamic decision making. Organizational Behavior and Human Decision Processes 43, 301–335.

Strohschneider, S., 1986. Zur Stabilität und Validität von Handeln in komplexen Realitätsbereichen. Sprache and Kognition 5, 42–48.

Stäudel, T., 1987. Problemlösen, Emotionen und Kompetenz. Die Überprüfung eines integrativen Konstrukts. Regensburg: Roderer.

Toda, M., 1962. The design of a fungus-eater: A model of human behavior in an unsophisticated environment. Behavioral Science 7, 164–183.

Wearing, A., C. Pivec and M.M. Omodei, 1991. Modeling strategies in dynamic decision making. Paper presented at the Psychonomic Society Meetings, San Francisco, November.